

HP LeftHand SAN Solutions

Support Document

User Manuals

Multi Site SAN User Manual for SANiQ 8.0



Legal Notices

Warranty

The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

The information contained herein is subject to change without notice.

Restricted Rights Legend

Confidential computer software. Valid license from HP required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Copyright Notices

© Copyright 2009 Hewlett-Packard Development Company, L.P.

Documentation Updates

The title page of this document contains the following identifying information:

- Software Version number, which indicates the software version.
- Document Release Date, which changes each time the document is updated.
- Software Release Date, which indicates the release date of this version of the software.

To check for recent updates or to verify that you are using the most recent edition of a document, go to:

<http://h20230.www2.hp.com/selfsolve/manuals>

This site requires that you register for an HP Passport and sign in. To register for an HP Passport ID, go to:

<http://h20229.www2.hp.com/passport-registration.html>

Or click the **New users - please register** link on the HP Passport login page.

You will also receive updated or new editions if you subscribe to the appropriate product support service. Contact your HP sales representative for details.

The following table indicates changes made to this document since the last released edition.

Document Changes

Chapter	Version	Changes

Support

Visit the HP Software Support web site at:

www.hp.com/go/hpsoftwaresupport

This web site provides contact information and details about the products, services, and support that HP Software offers.

HP Software online support provides customer self-solve capabilities. It provides a fast and efficient way to access interactive technical support tools needed to manage your business. As a valued support customer, you can benefit by using the support web site to:

- Search for knowledge documents of interest
- Submit and track support cases and enhancement requests
- Download software patches
- Manage support contracts
- Look up HP support contacts
- Review information about available services
- Enter into discussions with other software customers
- Research and register for software training

Most of the support areas require that you register as an HP Passport user and sign in. Many also require a support contract. To register for an HP Passport ID, go to:

<http://h20229.www2.hp.com/passport-registration.html>

To find more information about access levels, go to:

http://h20230.www2.hp.com/new_access_levels.jsp

Contents

1	Chapter: Designing A Multi-Site SAN.....	7
	Overview.....	7
	Concepts and Terms.....	9
	Designing A Multi-Site SAN.....	12
	Setting up Application Clusters in a Multi-Site SAN Environment.....	29
2	Chapter: Implementing A Multi-Site SAN.....	35
	Implementing the Multi-Site SAN.....	35
	Installing the Multi-Site SAN.....	36
	Converting an Existing SAN to a Multi-Site SAN.....	46
	Configuring Managers.....	55
	Sites Overview.....	58
	Adding Storage Nodes to a Multi-Site Cluster.....	64
	Removing Storage Nodes from a Multi-Site Cluster.....	66
	Index.....	69

This page is intentionally left blank

1 Chapter: Designing A Multi-Site SAN

Overview

The Multi-Site SAN features enable you to synchronously mirror data between sites automatically by configuring equal numbers of storage nodes into sites in the software. Capabilities enabled by the Multi-Site SAN software include:

- Multi-Site clusters that can
 - span up to 3 sites
 - span multiple subnets
 - have multiple VIPs
- Geographical awareness by designating storage nodes as members of a particular site
- Synchronously replicated (mirrored) data between sites based on volume level replication.
- The SAN/iQ[®] software uses the site information to ensure that data is mirrored between sites for high availability and disaster recovery.
- I/O path preferencing so that application servers connect to storage nodes located in the same site.
- Replication level 4 (four mirrors) of volumes that span the Multi-Site cluster.
- Failover Manager support for automatic failover/failback in 3-site configurations without requiring a physical storage node in the 3rd site.
- Failover Manager for quorum management if the network connection between the two sites becomes unavailable.
- Recover quorum via CLI for such time when a site failure also results in losing quorum in the management group.

Failover Manager Overview

The Failover Manager is a specialized version of the SAN/iQ software designed to run as a virtual appliance in a VMware environment. The Failover Manager participates in the management group as a real manager in the system; however, it does quorum operations only, not data movement operations. It is especially useful in a Multi-Site SAN configuration to manage quorum for the multi-site configuration without requiring additional physical hardware in a site.

The Failover Manager is supported on VMware Server, Workstation and Player. Find detailed information about installing and configuring the Failover Manager in [“Installing and Configuring a Failover Manager” on page 315 of Chapter 10, “Using Specialized Managers” in the *LeftHand SAN User Manual*](#).

Requirements

- A Multi-Site SAN requires a feature key for each storage node in the management group that is also participating in a Multi-Site cluster.
- All sites must have the same number of storage nodes. You must add storage nodes to, or remove them from, Multi-Site clusters equally.

Concepts and Terms

Table 1 Multi-Site SAN Glossary (Continued)

Term	Definition
Site	A user-designated location in which storage nodes are installed. Multi-Site SAN configurations have multiple sites with storage nodes in each site, and each site has its own subnet. A site can be a logical configuration, such as a subnet within the same data center, department, or application.
Primary Site	<p>A site designation assigned by the administrator in the Centralized Management Console. A primary site is more important than a secondary site. In this setup, you would run a majority of managers in the primary site. In a two-site setup, this allows the primary site to stay online even if the network link between the primary and secondary sites fails. Typically, the primary site has majority/all of the application servers.</p> <p>In configurations that do not designate a primary site, the sites are referred to as “peer” sites.</p>
Secondary Site	A site that is less important than the primary site. In this setup a minority of managers runs in the secondary site. In a two-site setup, this allows the secondary site to go offline if the network link between the Primary and secondary sites fails. Typically, the secondary site has a minority, or none, of the application servers. If the primary site fails, customers can manually recover quorum in the secondary site.
Logical Site	This site is on an isolated network and power connection than the other sites. However, it can be in the same physical location as one of the real sites. Also, a site for a Failover Manager.
Peer Site	Absence of a Primary site designation makes all the sites peer sites.

Table 1 Multi-Site SAN Glossary (Continued)

Term	Definition
Disaster Recovery Site	Similar to a secondary site, the disaster recovery site is used to operate the SAN in the event of a disaster.
Data Center	Also known as a “Site.” A data center is a physical location in your environment where application servers, SAN storage and network equipment reside. In the SAN/iQ Multi-Site software, a data center is typically referred to as a site.
Cluster	A cluster is a grouping of storage nodes that create the storage pool from which you create volumes.
Standard Cluster	Also known as a “cluster.” A standard cluster is one that does not use any of the Multi-Site features within the SAN/iQ software. Standard clusters <ul style="list-style-type: none">• Cannot contain storage nodes that are designated to reside in a site.• Cannot contain storage nodes that span subnets• Can only have a single VIP.
Manager	Managers are storage nodes within a management group that you designate to govern the activity of all of the storage nodes in the group.
Multi-Site Cluster	A cluster of storage that spans multiple sites (up to 3). A Multi-Site cluster must meet at least one of the following conditions: <ul style="list-style-type: none">• Contain storage nodes that reside in two or more sites• Contain storage nodes that span subnets• Contain multiple VIPs. The cluster can have a single site, and the multiple VIPs make it a multi-site cluster.

Table 1 Multi-Site SAN Glossary (Continued)

Term	Definition
Failover Manager	A specialized manager running as a VMware appliance that allows you to place a quorum tie-breaker node into a 3rd location in the network to provide for automated failover/failback of the Multi-Site SAN clusters. The Failover Manager is designed to run on VMware ESX Server, VMware Server and VMware Player. It is installed on hardware separate from the SAN hardware.
Virtual Manager	A manager that is added to a management group but is not started on a storage node until it is needed to regain quorum.
Quorum	A majority of managers required to be running and communicating with each other in order for the SAN/iQ software to function.
Volume	A logical entity that is made up of storage on one or more storage nodes. It can be used as raw data storage or it can be formatted with a file system and used by a host or file server.
Replication	Designate how many copies of data you want to keep in the cluster.
Network RAID	Designate how many copies of data you want to keep in the cluster.
Network RAID	Also known as “replication.” Synchronous replication, or mirroring on a volume-by-volume basis. Mirroring of data for a volume across all storage nodes in the cluster. Network RAID level 2 or higher is required to replicate data in a Multi-Site SAN environment.

Designing A Multi-Site SAN

The Multi-Site SAN software offers multiple features and the design flexibility to protect against certain types of failures in the environment.

Data replication, site design and quorum managers combine to protect against certain types of failures in the environment. Some of these common types of failures a Multi-Site SAN protects against include:

Data Center Failures

- Site power outage
- Site network outage
- Site disaster (fire, flood, terrorist attack, etc.)

Individual Storage Node Failures

- Power outage
- Network outage
- Node disaster
- System failure (backplane, motherboard, RAM, CPU)

Configuring Data Replication for Optimal High Availability and Redundancy

The SAN/iQ software offers four levels of synchronous replication at a volume level. In order to protect the data across sites, you must choose a replication level that is at least equal to the number of sites in the configuration. For instance, if the storage cluster spans 3 sites, then volumes must be 3-way replicated in order to insure that data is the same on each site.

[Table 2](#) and [Table 3](#) summarize the data replication levels for the Multi-Site SAN and associated supported configurations.

Table 2 Volume Replication Levels and the Supported Number of Sites

Volume Replication Level	Supported Number of Sites
2	2
3	2 or 3
4	2 or 3

Table 3 Number of Sites and the Supported Volume Replica

Number of Sites	Supported Replication Levels
2	2,3, or 4
3	3, 4

When to Use 4-Way Replication

Using 4-way replication offers a level of data protection that enables you to sustain a complete site failure and still maintain replicated data in the remaining site. For instance, with 4-way replication there are four copies of data for the volume on the SAN. In a 2-site configuration, the Multi-Site SAN software puts 2 of the 4 copies in one site and the remaining 2 in the other. If a site fails, the remaining site still has 2 copies and remains faulttolerant.

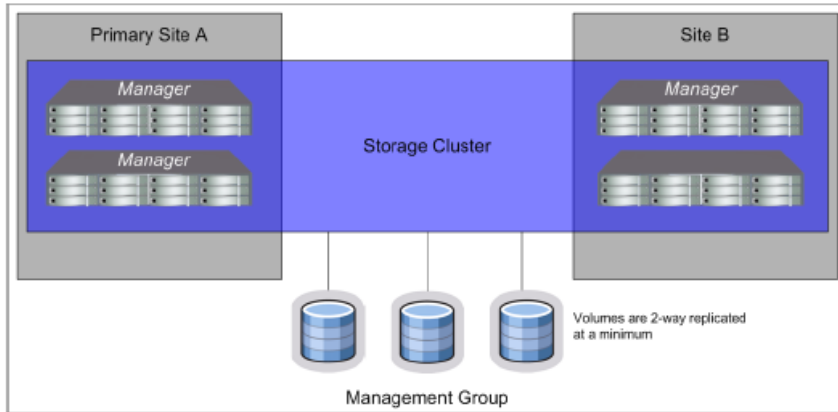


Figure 1 Sample configuration for 2-site, 2-way replication

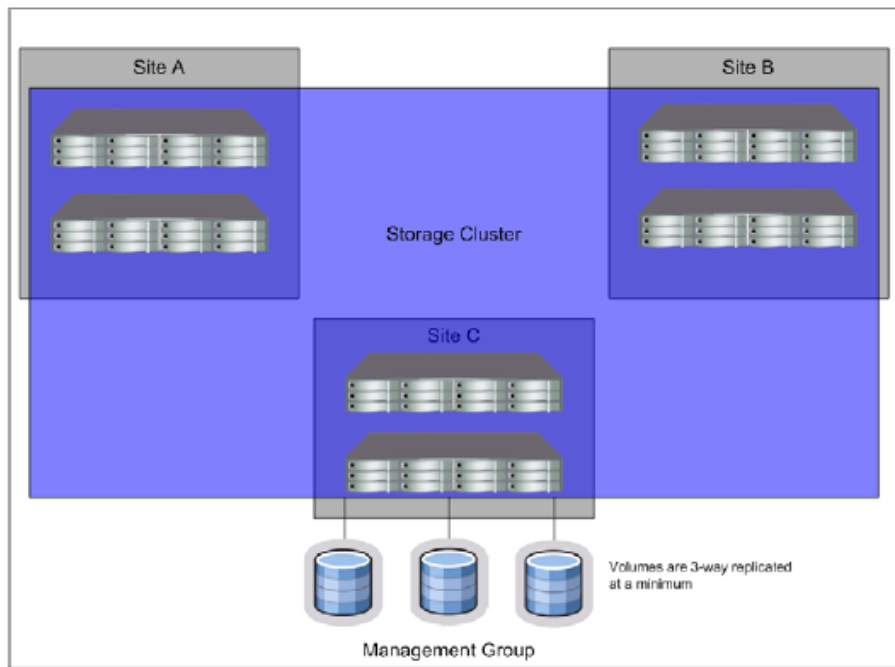


Figure 2 Sample Configuration for 3-site, 3-way replication

Configuring Sites and Quorum Managers for Optimal High Availability and Redundancy

There are several recommended site configurations that can be used with the Multi-Site SAN software that provide varying levels of availability and redundancy. The following table summarizes the recommended configurations. Select the one that most appropriately matches your environment.

In addition to setting up the Multi-Site cluster(s) for data replication, it is important to correctly set up the SAN managers for quorum, to ensure that in the event of a failure, the system can be properly recovered. Maintaining quorum can be either an automated process or manually recovered by an Administrator. The following table summarizes the common Multi-Site SAN and manager configurations:

Table 4 Common Configurations of Multi-Site SANs and Managers

Number of Sites	Cluster/Site Setup	Total Number of Managers ^a	Failover Manager	Site Types
3 physical sites	3 physical sites	Min. = 3 Max = 5	In the 3 rd -site	In the 3 rd -site
2 physical sites, with logical 3 rd site	Cluster spans 2 sites	Min. = 3 Max = 5	In the logical 3 rd -site	All sites are peer
3 physical sites	Cluster spans 3 sites	Min. = 3 Max = 6	Not Used	All sites are peer
2 physical sites	Cluster spans 2 sites	Min. = 3 Max = 5	Not Used	One site is designated Primary

^aIncludes all manager types: Manager, Failover Manager, and Virtual Manager.

Best Practice

In a cluster that spans multiple sites, run the same number of managers in each site. In a cluster that spans 2 sites, run a Failover Manager in a 3rd site (physical or logical) that is not included in the cluster. The one exception to this rule is the 2-site configuration where one site is configured as primary. In this case, run a majority of managers in the primary site.

Sample Configurations

The following graphics depict common sample configurations of the cases mentioned above. Note that management groups and Multi-Site clusters can have many more storage nodes than is depicted here. However, the graphics layouts illustrate the significant points of site-to-cluster mappings and the appropriate placement of managers.

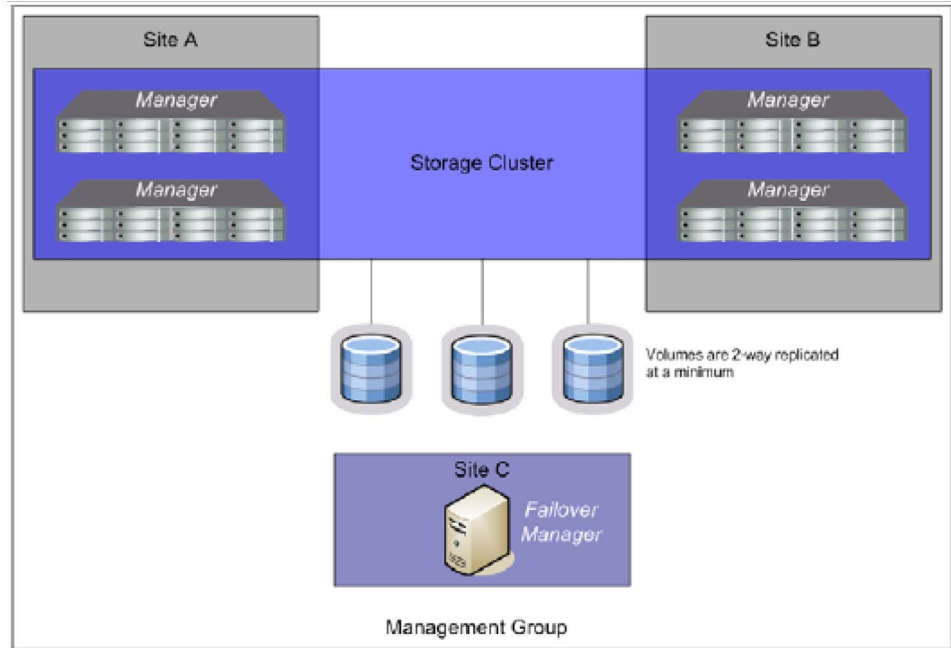


Figure 3 3 physical sites, cluster spans 2 sites, Failover Manager in 3rd site

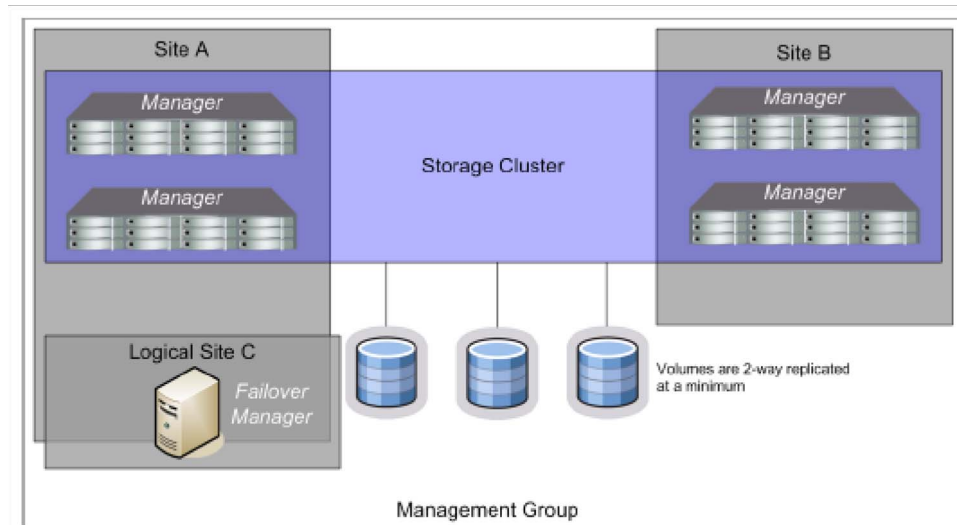


Figure 4 physical sites, logical 3rd site, cluster spans 2 sites, Failover Manager in logical 3rd site

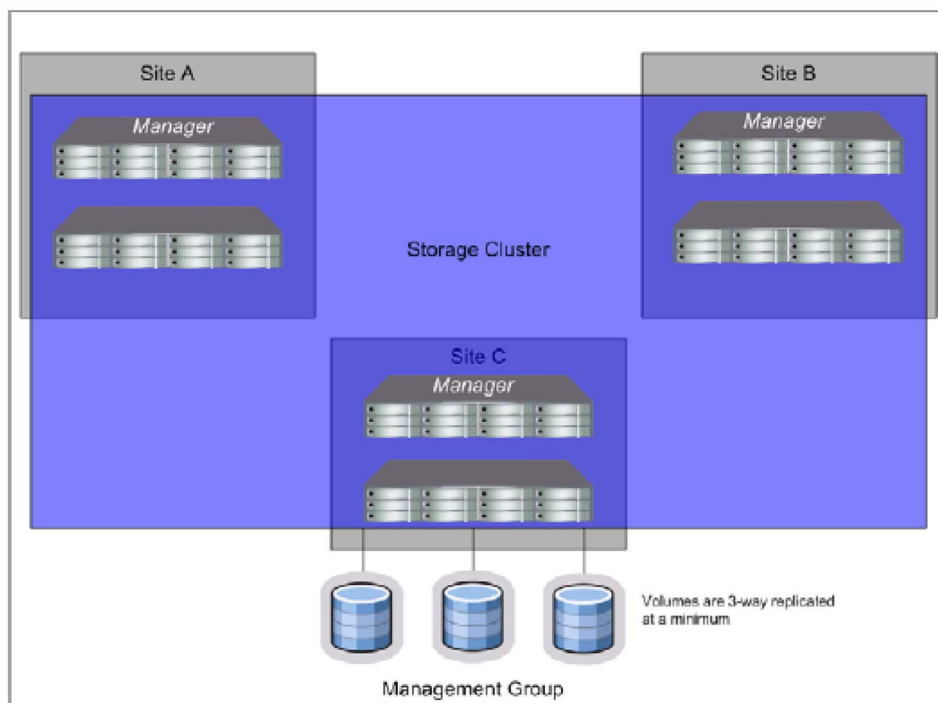


Figure 5 3 physical sites, cluster spans 3 sites

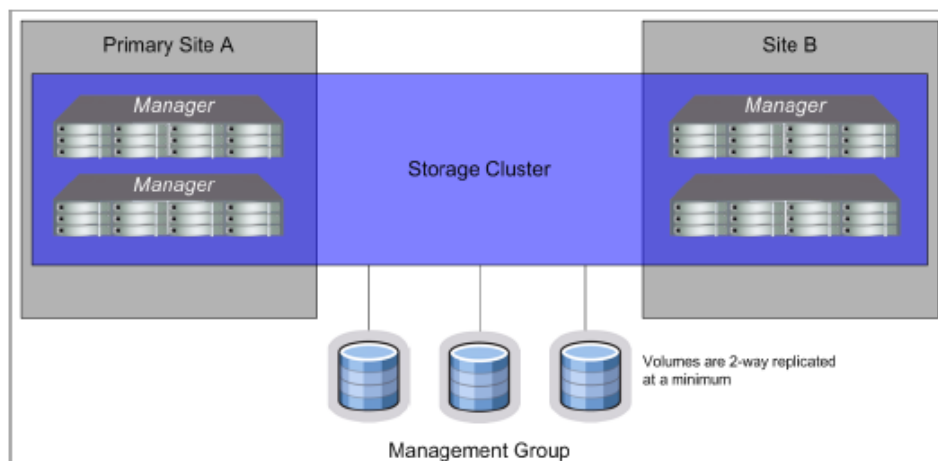


Figure 6 2 physical sites, cluster spans 2 sites, one site is the primary site

Alternative Setups

The Multi-Site SAN software supports many other valid configurations than the four mentioned in the prior section. [Table 5](#) lists other possible configurations.

Table 5 Alternative Multi-Site SAN Configurations

Alternate Configuration	Description	Details
More than one Multi-Site cluster in the management group	Within the same management group, you run multiple clusters.	The Multi-Site SAN software supports running many clusters in the same management group. You simply need to make sure that clusters have the same number of storage nodes in each site for a particular cluster.
More than 3 sites	Within the same cluster, you have more than 3 sites.	A particular cluster can only span up to 3 sites. However, the Multi-Site SAN software supports having more than one cluster. The different clusters can span different sites within the same cluster.

Table 5 Alternative Multi-Site SAN Configurations

Alternate Configuration	Description	Details
Using Virtual Manager	The Virtual Manager capability remains in the product. Virtual Manager is a manual process for recovering quorum.	The Virtual Manager is used in 2- site configurations, and is considered a legacy implementation. See Chapter 12, “Disaster Recovery Using a Virtual Manager” in the <i>LeftHand SAN User Manual</i> .
Remote Copy with Multi-Site SAN	Remote Copy and all other SAN/iQ features are completely compatible with Multi-Site SAN.	You can use Remote Copy to replicate data from a Multi-Site SAN to a Disaster Recovery or Backup site. See Chapter 17, “Understanding and Planning Remote Copy” in the <i>Remote Copy User Manual</i> .
Multiple Subnet clusters in the same site	Customers can use the Multi-Site SAN features to set up clusters that span multiple subnets in the same site.	The Multi-Site SAN software can be used to create clusters that span multiple networks (subnets) in larger Enterprise data centers. The storage nodes all reside in the same site. However, the cluster must be Multi-Site if it is to contain storage nodes from different subnets.

Designing the Network for the Multi-Site SAN

Good network design is a critical part of setting up a Multi-Site SAN that meets your needs. A good Multi-Site SAN network encompasses the following attributes:

- Reliable
- Highly Available
- High Performance

There are a set of minimum recommended network design considerations when planning out the network that supports the Multi-Site SAN.

Best Practices for Implementing the Multi-Site SAN Network

- Adequate bandwidth — Plan for 50MB/sec of bandwidth for each storage node in each site. For instance, if each site contains 5 storage nodes, then you need 250 MB/sec throughput. In this case, that translates into two Gigabit Ethernet links, or more.
- Low round-trip latency — In order to not impact disk I/O to the application server, the round-trip latency between the sites must be no more than 2ms.
- Redundant links — Between the sites, you should have multiple physical connections (media) between the sites for redundancy. The network should be configured so that a link failure does not cause the Multi-Site SAN to go offline.
- Use Multiple Subnets — Each site should contain its own subnet. When the Multi-Site cluster spans different subnets, other features can be used:
 - Optimal I/O Pathing = application servers will perform I/O only against storage nodes that are in the same site (subnet) as the server.
 - Multiple VIPs = a VIP is assigned to each site/subnet pair. A cluster can have more than one VIP.
- Use the Failover Manager — Unless you have storage nodes in 3 physical sites, use a Failover Manager. Put the Failover Manager in the 3rd site (physical or logical). A 3rd site for Failover Manager should be on an isolated network and power connection than the other sites.

Common Multi-Site SAN Network Designs

Based on the Multi-Site SAN configuration, there are many network designs that can be implemented that solve the Redundancy, High Availability and Performance requirements for the Multi-Site SAN. The following are common network designs:

- Dual Redundant Links between the sites
- Full-Mesh Triangular (3) Redundant Links between sites
- Hub & Spoke Central network core with redundant links out to each site
- Combination Full-Mesh Core network w/ Triangular (3) links between sites

Balancing your requirements for High Availability during certain types of failures in the Multi-Site SAN and the budget associated with the network infrastructure typically dictates the network architecture a customer will choose to implement. Obviously, the most redundant network typically will cost the most to implement. However, not implementing redundancy in the Multi-Site SAN network will significantly curtail the benefits of the Multi-Site SAN software. The following table summarizes the common Multi-Site network configurations and the types of failures for which they protect:

Table 6 Common Multi-Site Network Configurations and Failure Protection of Each

Network Topology	Site Failure	Network Failure	Storage Node Failure
Dual Redundant Links between sites	Yes, with manual intervention	Yes	Yes

Table 6 Common Multi-Site Network Configurations and Failure Protection of Each

Network Topology	Site Failure	Network Failure	Storage Node Failure
Full-Mesh Triangular Links between sites	Yes	Yes	Yes
Hub & Spoke central core	Yes	Yes, so long as the Core does not fail	Yes
Combo Full-Mesh Core w/ Triangular links between sites	Yes	Yes	Yes

Sample Recommended Network Configurations for Multi- Site SANs

The following graphics depict common sample network configurations of the cases mentioned above. The clusters and sites can have many more storage nodes than is depicted here, however, the pictures are more focused on the network design.

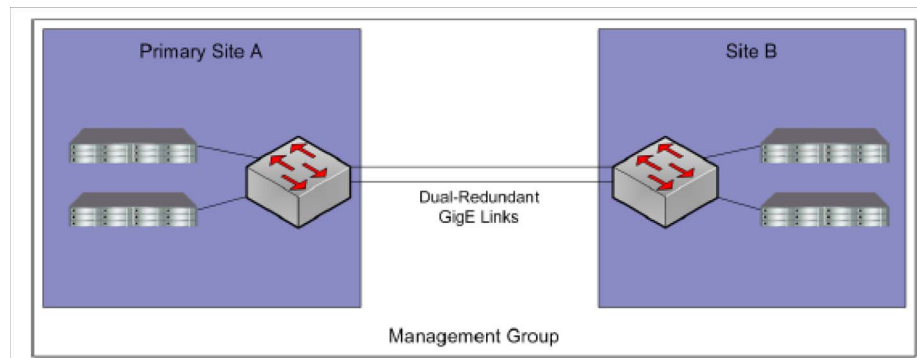


Figure 7 Dual Redundant Links between 2 sites

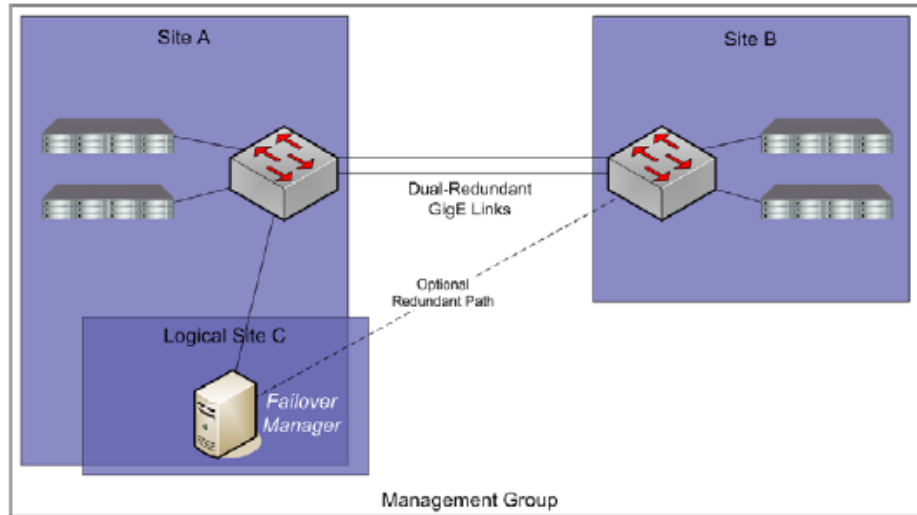


Figure 8 Dual redundant links between 2 sites with Failover Manager in logical 3rd site

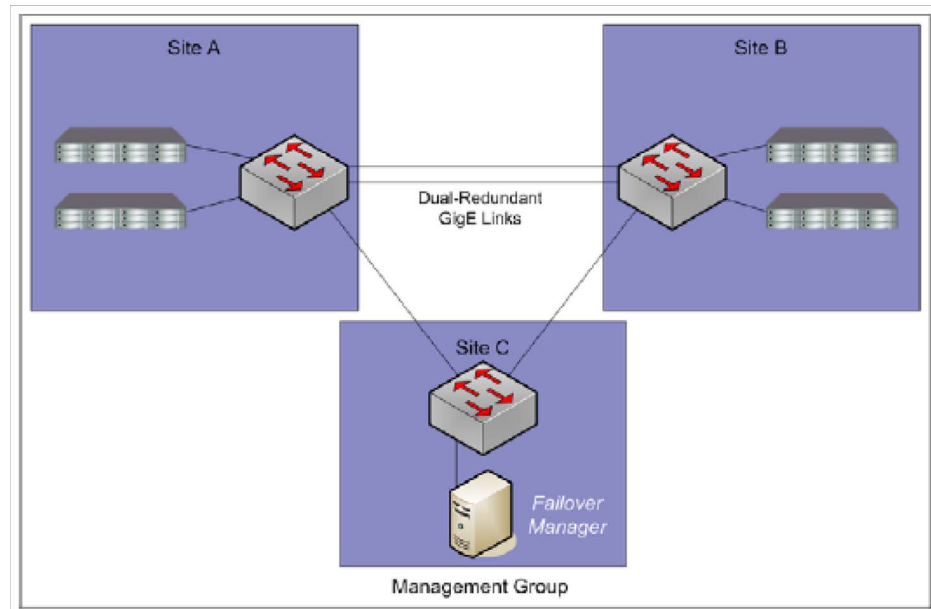


Figure 9 Triangular Network with Failover Manager

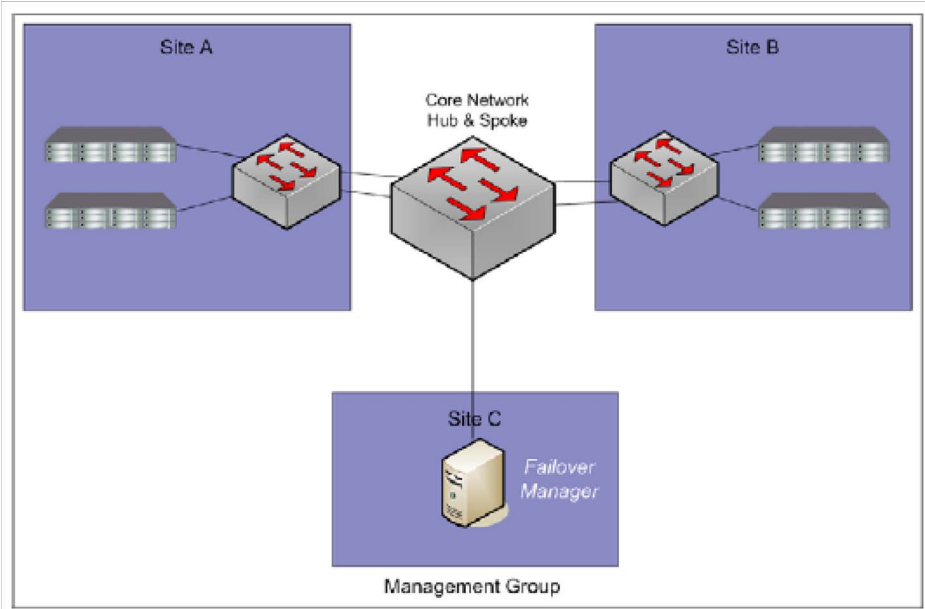


Figure 10 Hub and spoke network core, redundant links out to each site

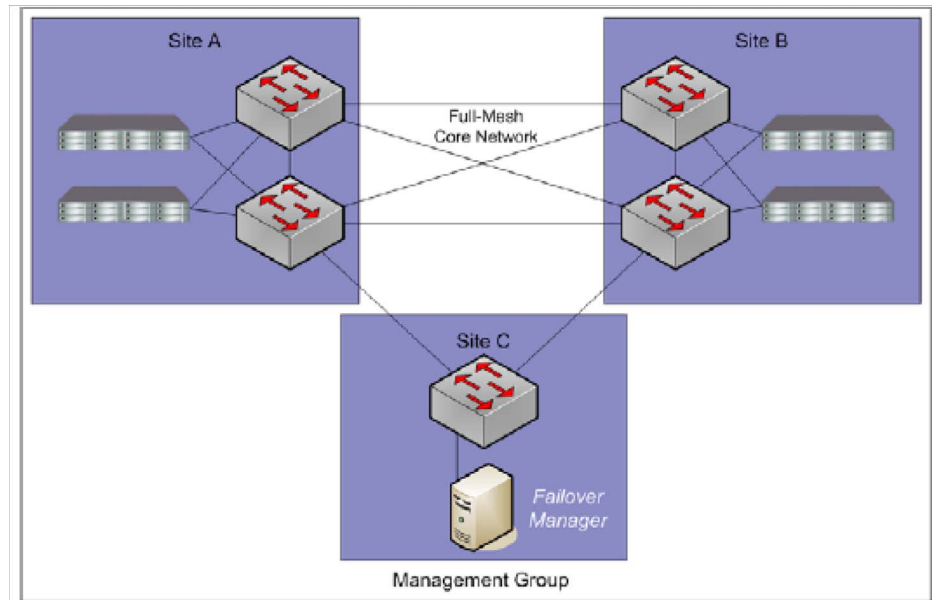


Figure 11 Combination full-mesh core with triangular links to all sites

Setting up Multiple Subnets for the Multi-Site SAN

Many of the advanced features in the Multi-Site SAN software depend on the implementation of multiple subnets within the iSCSI cluster. Preferably, customers should implement a subnet per site. By doing so, you can take advantage of the following Multi-Site SAN features:

- Virtual IP Address (VIP) per site. iSCSI discovery sessions are terminated on a storage node that is local to the application server in that site.
- Virtual IP Addresses (VIPs) are only hosted on storage nodes that reside in the same subnet as the VIP.
- Read and Write I/O requests are forwarded from the application server to a storage node that is in the same site as the server. This ensures that the I/O path is optimal and multiple network hops across the network link between the sites are eliminated.

The following graphic is an example of how the Multi-Site SAN setup maps to subnets, VIPs and application servers.

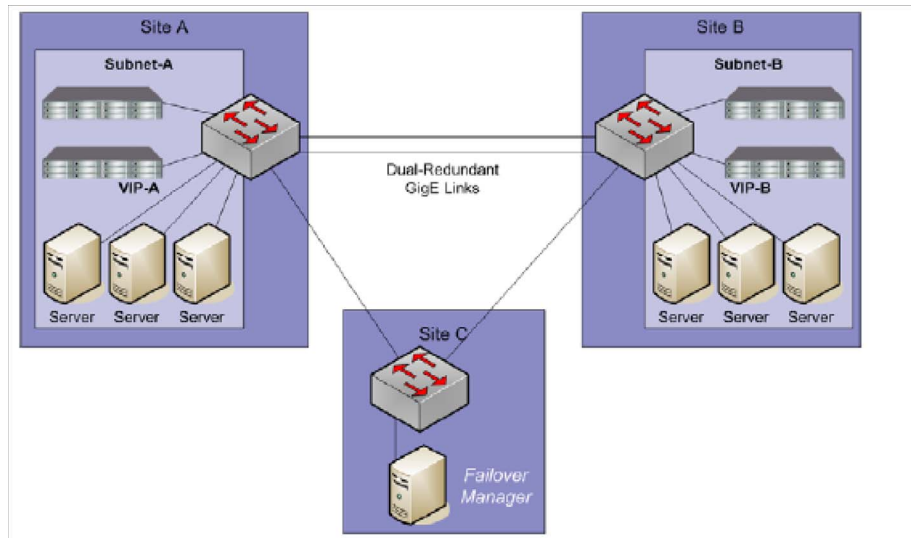


Figure 12 Multi-Site SAN mapping to subnets, VIPs and application servers

Using the DSM for MPIO with Multiple Subnets

When you are using multiple subnets in a multi-site SAN, the DSM for MPIO will only build I/O paths to the storage nodes that are in the same subnet as the application server. See the SAN/iQ[®] Solution Pack for Microsoft[®] Windows[®] *User Guide for more information about DSM for MPIO.*

Setting up Application Clusters in a Multi-Site SAN Environment

Similar to how you can set up a Multi-Site SAN for data availability even when a complete site is offline, application clusters (Microsoft Cluster Server, VMware ESX Cluster, Linux Cluster, etc.) can be set up so that the cluster nodes are split between the sites. By having the automated failover/failback capabilities in the SAN, many application cluster technologies can leverage this capability to do “stretch” application clustering. Consult the application server specific information on how to set up the clustering to support this type of failover. At a high level, here are the general steps:

- 1 Set up an application server node in each site.
- 2 Configure the application server to connect to the volumes on the SAN.
- 3 If using multiple subnets in the storage cluster, log on to both VIPs in the cluster.

If you are using the LeftHand Networks® DSM for MPIO, refer to the chapter, “Using the DSM for MPIO,” in the *SAN/iQ® Solution Pack for Microsoft® Windows® User Guide*.

Using the Primary Site Designation in 2-Site Configurations

The Multi-Site SAN software supports designating a particular site as Primary so that it does not go offline when the secondary site goes offline, or when the network link between sites goes offline. This is specific to the 2-site configuration where, without this designation, there would be no way for the system to determine which site should remain online in the event of a network separation, or what is classically known in the clustering/ distributed systems world as a “split-brain” system. To set up this type of configuration, you simply do the following:

- 1 Designate one of the sites as primary.
- 2 Start a majority of managers in the primary site.

Failure Modes and Recovery Steps

In general, there are three failure modes that a configuration like this can encounter that you should be aware of. The following table summarizes the failure conditions and what you should do to rectify the issue.

Table 7 Failure Conditions and Resolutions

Failure Condition	Resolution Steps
Secondary site offline	None. The primary site will continue to operate.
Network Link between sites fails	None. The primary site will continue to operate.
Primary site offline	<p>If any of the storage nodes running managers in the primary site can be brought back online, you should do this first.</p> <p>If all storage nodes in the primary site are lost forever, you can use the Recover Quorum command-line function to recover quorum in the secondary site manually.</p>

Recover Quorum Operation

The Recover Quorum function is introduced with the Multi-Site SAN as a way for you to bring volumes back online in the event that quorum is lost in the management group.



Warning: The Recover Quorum function should only be used as a last resort. The failed site must be completely failed to use Recover Quorum safely on the remaining site.

The Recover Quorum command-line function does the following on the storage nodes in the management group:

- 1 All offline managers are removed from the management group.
- 2 All Offline storage nodes are ejected from the management group.

- 3 The Global Database for the management group is incremented by one billion operations, sufficient to prevent any data integrity issues in the event any of the offline storage nodes are brought back online.

Best Practices for Quorum and Manager Configuration in a Multi-Site SAN Environment

Configuring the appropriate number of managers in the Multi-Site SAN on the appropriate storage nodes is an important part of setting up the Multi-Site SAN for high availability and fault-tolerance. The following table summarizes common Multi-Site SAN configurations and the recommended manager configurations:

Table 8 Common Multi-Site SAN Configurations with Recommended Managers

Site Setup	Manager Configuration	Quorum Count
2 physical sites, one site is primary	<ul style="list-style-type: none">• 4 storage nodes: run 2 managers in the primary site, 1 manager in the secondary site• 6 storage nodes or more: run 3 managers in the primary site, 2 managers in the secondary site	<ul style="list-style-type: none">• Manager Count = 3, Quorum Count = 2• Manager Count = 5, Quorum Count = 3

Table 8 Common Multi-Site SAN Configurations with Recommended Managers

Site Setup	Manager Configuration	Quorum Count
2 physical sites, logical 3 rd site	<ul style="list-style-type: none"> • 2 storage nodes: run 1 manager in each site, 1 failover manager in the logical 3rd site • 4 storage nodes or more: run 2 managers in each site, 1 failover manager 	<ul style="list-style-type: none"> • Manager Count = 3, Quorum Count = 2 • Manager Count = 5, Quorum Count = 3
3 physical sites, cluster spans 2 sites	<ul style="list-style-type: none"> • 2 storage nodes: run 1 manager in each site, 1 failover manager in the 3rd site • 4 storage nodes or more: run 2 managers in each site, 1 failover manager in the 3rd site 	<ul style="list-style-type: none"> • Manager Count = 3, Quorum Count = 2 • Manager Count = 5, Quorum Count = 3
3 physical sites	3 storage nodes: run a manager in each site 6 or more storage nodes: run 2 managers in each site	Manager Count = 3, Quorum Count = 2 Manager Count = 6, Quorum Count = 4

2 Chapter: Implementing A Multi-Site SAN

Implementing the Multi-Site SAN

Once you have designed the Multi-Site SAN configuration, you are ready to implement it.

Before You Begin

Install Multi-Site SAN feature keys on the storage nodes you are going to use.

See “Registering Add-on Applications” on page 451 in Chapter 17, “Feature Registration.” of the *LeftHand SAN User Manual* for information about registering your storage nodes.

Have the following information ready.

- The layout for storage nodes and sites
- IP addresses, subnets and VIPs to use for the Multi-Site cluster. You can use one VIP per subnet.
- (Optional) IP address and subnet for the Failover Manager

Preparing to Install a New Multi-Site SAN

Once you have finished planning the Multi-Site SAN and installed feature keys on the storage nodes for the SAN, you are ready to begin installation.

Preparing to Convert an Existing SAN to a Multi-Site SAN

Converting an existing SAN to a Multi-Site SAN includes planning the site layout and network information along with the following considerations:

- Do you need to move/remove data from the existing SAN before reconfiguring it?
- Do you need to purchase additional storage nodes to meet the requirement that each site have the same number of storage nodes?
- Do you need to reconfigure your network to take advantage of the Multi-Site SAN capabilities?

See [Designing the Network for the Multi-Site SAN](#) on page 22 for detailed information about setting up your network appropriately.

Installing the Multi-Site SAN

To install a Multi-Site SAN, use the Management Groups, Clusters, and Volumes Wizard, shown in [Figure 13](#)

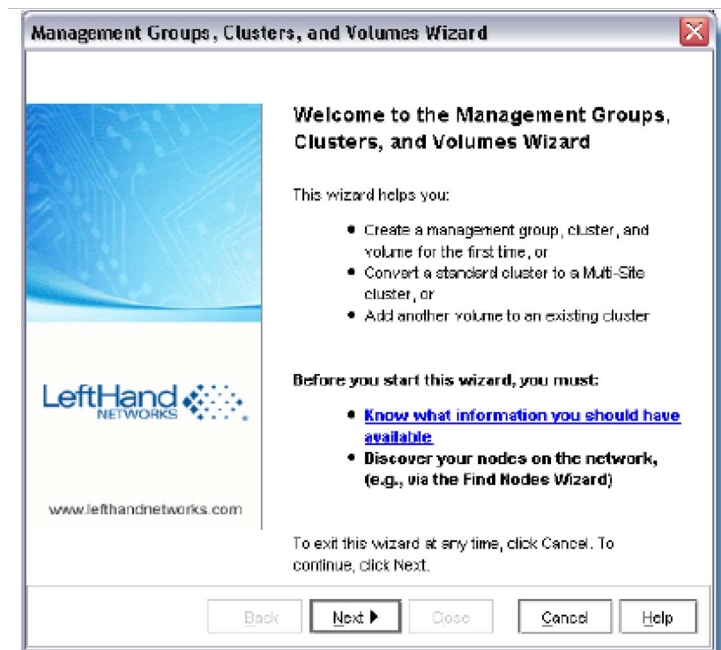


Figure 13 Creating a Multi-Site SAN with the Management Group Wizard

You must have at least one storage node for each site in the Available pool before starting.

Creating the Management Group

The first part of the wizard creates the management group, which includes configuring the optimal number of managers for the initial setup. After you have the Multi-Site clusters configured, you need to review the manager configuration and adjust it according to the Multi-Site configuration you have planned. See [Configuring Managers](#) on page 55

Creating Multi-Site Clusters and Volumes

The path through the Multi-Site cluster portion of the wizard follows these steps, described in greater detail in the next section, “[Using the Multi-Site Cluster Wizard](#)”.

- 1 Create a cluster.
- 2 Create a site and then assign storage nodes to the site.
- 3 Create the additional sites with storage nodes, up to 3 sites total.
- 4 Add one or more VIPs, depending on the number of subnets you have configured.
- 5 Create volumes using the replication levels dictated by your Multi-Site cluster design.

Using the Multi-Site Cluster Wizard

After creating the management group, you come to the Create a Cluster window of the wizard, shown in [Figure 14](#).

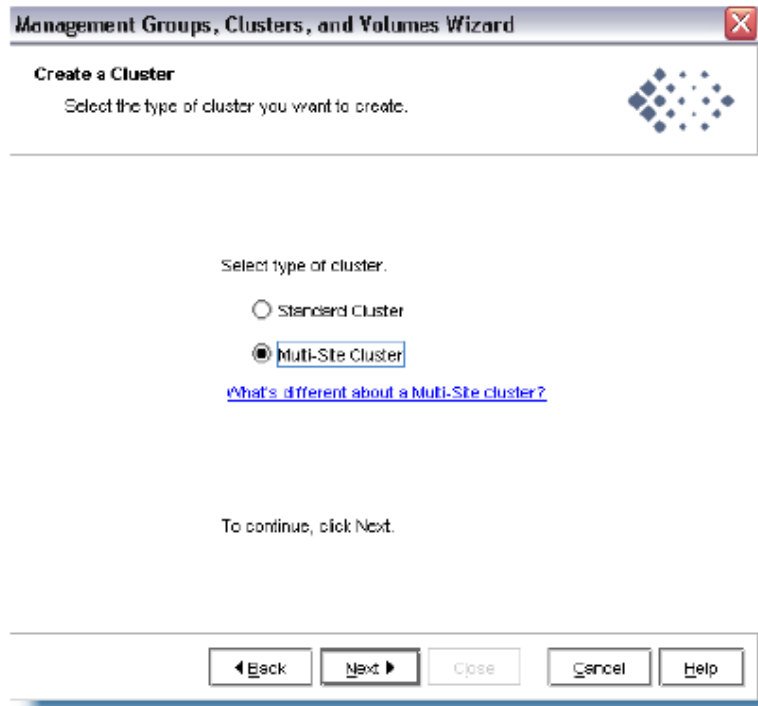


Figure 14 Creating Multi-Site clusters

- 1 Select Multi-Site Cluster and click Next.

The Create Cluster window opens, displaying all the storage nodes you designated for the management group, as shown in [Figure 15](#).

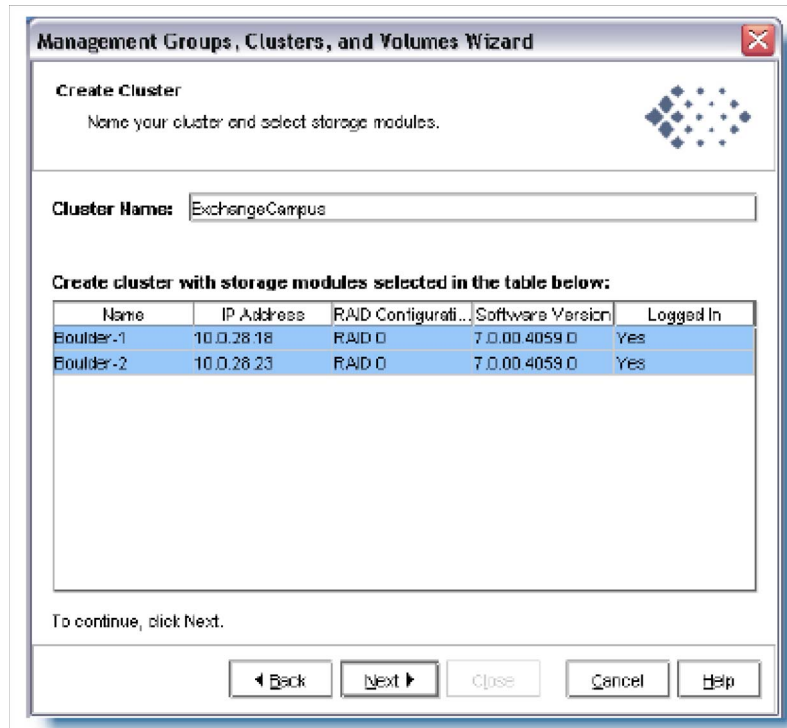


Figure 15 Selecting storage nodes for a Multi-Site cluster

- 2 Select all the storage nodes targeted for the all the sites you plan to create, type a name for the cluster, and click Next.

The Set up Sites window opens, shown in [Figure 16](#).

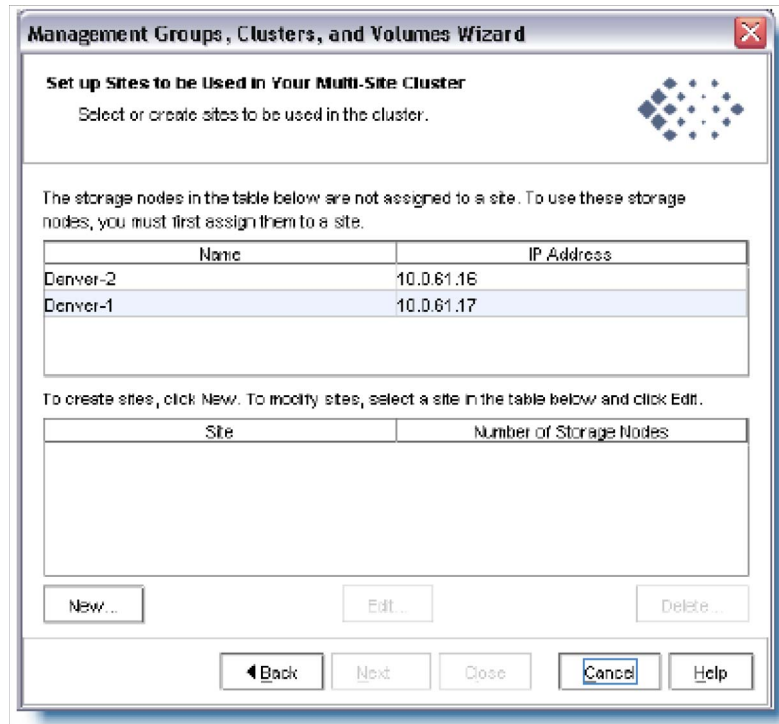


Figure 16 Setting up the sites for a Multi-Site cluster

Creating Sites

First you create a site and then you assign storage nodes to that site.

- 3 Click New to open the New Site window.
- 4 Enter a name for the site and an optional description.
- 5 [Optional] Make the site primary, if appropriate.
- 6 Click Add in the Site Nodes section to select the storage nodes for this site.

Figure 17 shows a completed New Site window.



Figure 17 Setting up a site and adding storage nodes

- 7 Click OK when you have completed creating the site. The Set up Sites window opens, with the first site created and the remaining storage nodes in the list, as shown in [Figure 18](#).

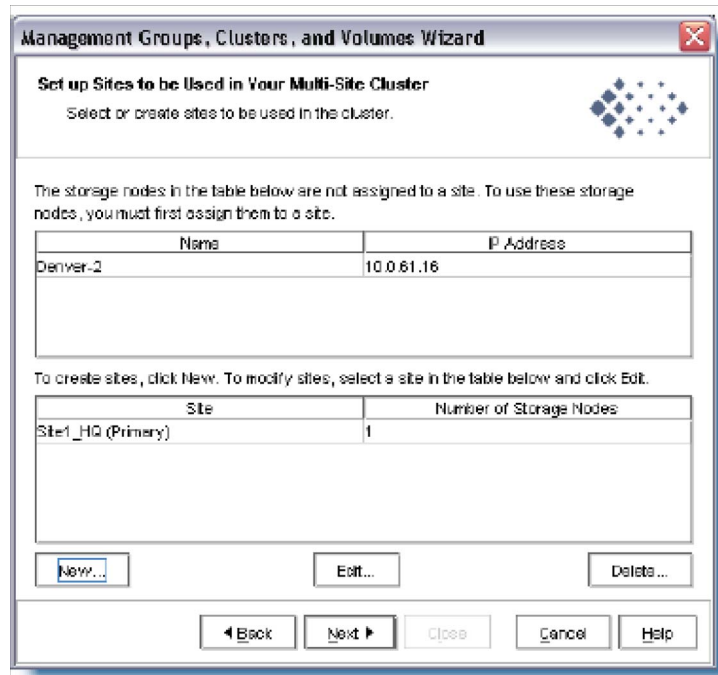


Figure 18 First site completed

- 8 Repeat [step 3](#) through [step 7](#) to create additional sites.
- 9 When you have finished creating sites, click Next to continue with the wizard.

The Assign VIPs and Subnets window opens, shown in [Figure 19](#).

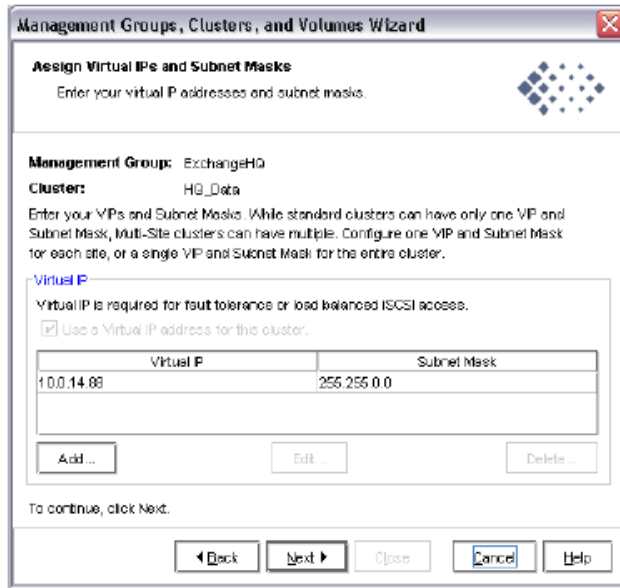


Figure 19 Adding VIPs and subnets to the Multi-Site cluster

- 10 Click Add to assign each VIP and Subnet Mask as planned for your site layout.

You can have 1 VIP per subnet.

- 11 Click Next to continue.

The Create Volume window opens, shown in [Figure 20](#).

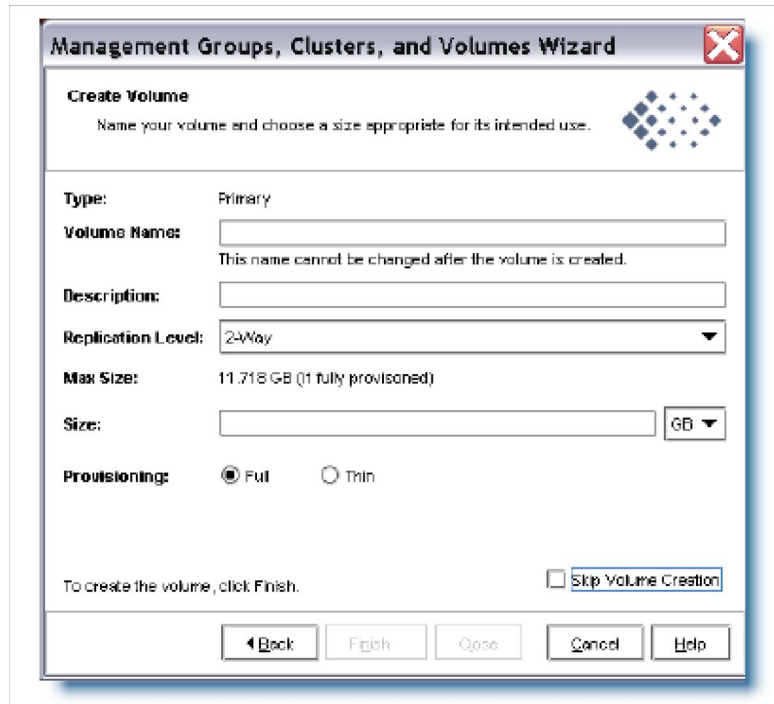


Figure 20 Adding a volume to the Multi-Site cluster

Completing the Wizard

- 12 Click Finish when you have completed defining the volume.

The management group, sites, Multi-Site cluster and the volume are all created. This process takes some minutes.

Viewing the Multi-Site Cluster Information

You can view the completed sites and related information in two places. The Cluster tab view shows the sites in relation to the cluster.

Cluster Tab View

Select the cluster in the navigation window. The cluster Details tab displays the sites, as shown in [Figure 21](#).

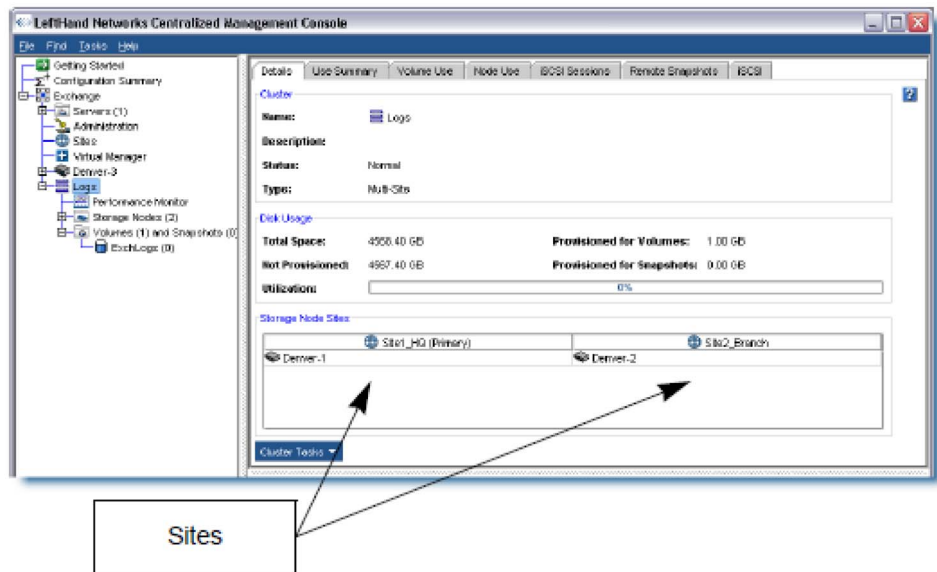


Figure 21 Viewing the sites in the cluster Details tab

In this site view, you can see the names of the sites and the storage nodes that reside in each.

Sites View

Select the Sites node in the navigation window to view detailed information about the storage nodes in the sites, as shown in [Figure 22](#).

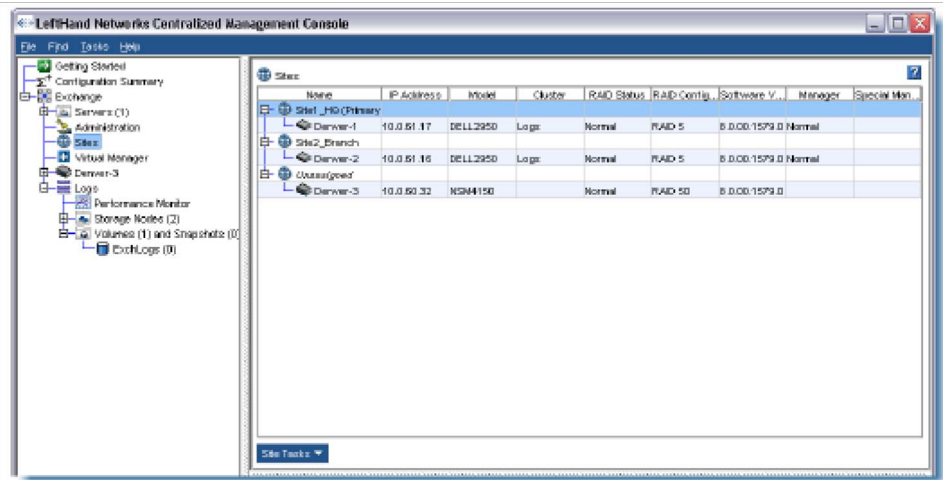


Figure 22 Viewing the Sites Details tab window

In this view of the sites, you can easily review details about the storage nodes, the cluster, and the management group.

Converting an Existing SAN to a Multi-Site SAN

Plan the conversion according to the criteria described in [Preparing to Convert an Existing SAN to a Multi-Site SAN](#) on page 35.

Prerequisites

- Upgrade all storage nodes to SAN/iQ software version 7.0 or later.
- Obtain the feature keys for storage nodes to be used in the Multi-Site cluster.
- Apply the feature keys to the storage nodes.
- Add or remove storage nodes to the cluster if necessary to ensure you can have equal numbers of storage nodes in each site. The wizard will not let you create sites with different numbers of storage nodes. Nor will it let you leave unassigned any storage nodes currently in the cluster.

Using the Multi-Site Cluster Wizard

Use the Management Groups, Clusters, and Volumes wizard to convert an existing standard cluster to a Multi-Site cluster.

- 1 Start the wizard from the Getting Started Launch Pad.
- 2 In the wizard, select Existing Management Group, as shown in [Figure 23](#), and click Next.

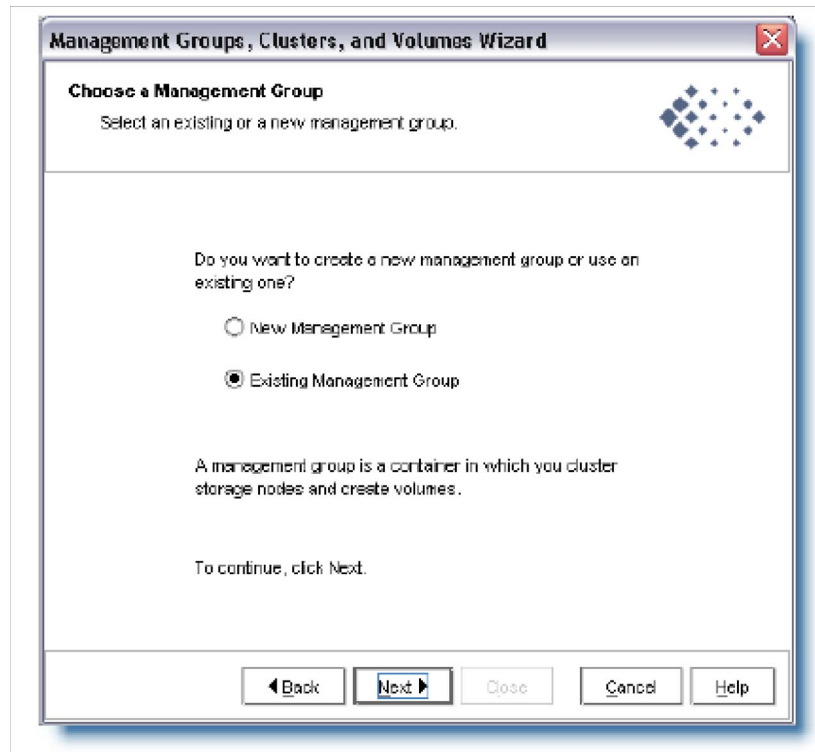


Figure 23 Begin the process by choosing an existing management group

- 3 Select the management group that contains the cluster you want to convert, as shown in [Figure 24](#).

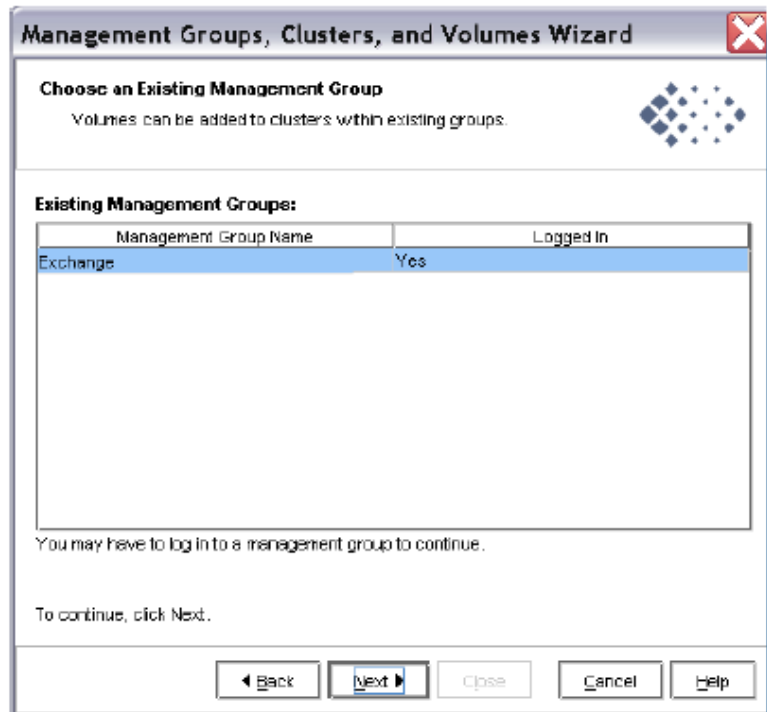


Figure 24 Select the specific management group

- 4 Click Next to open the Create a Cluster window.
- 5 Select Existing Cluster and Convert a Standard Cluster to a Multi-Site Cluster, as shown in [Figure 25](#), and click Next.

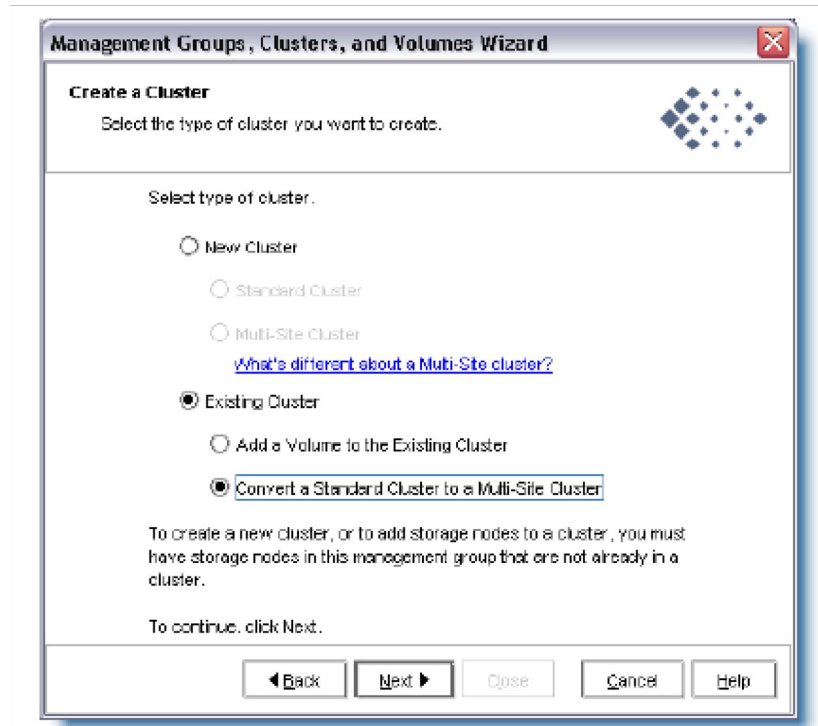


Figure 25 Preparing to convert a standard cluster to a Multi-Site cluster

- 6 Click Next to select the specific cluster, shown in [Figure 26](#).

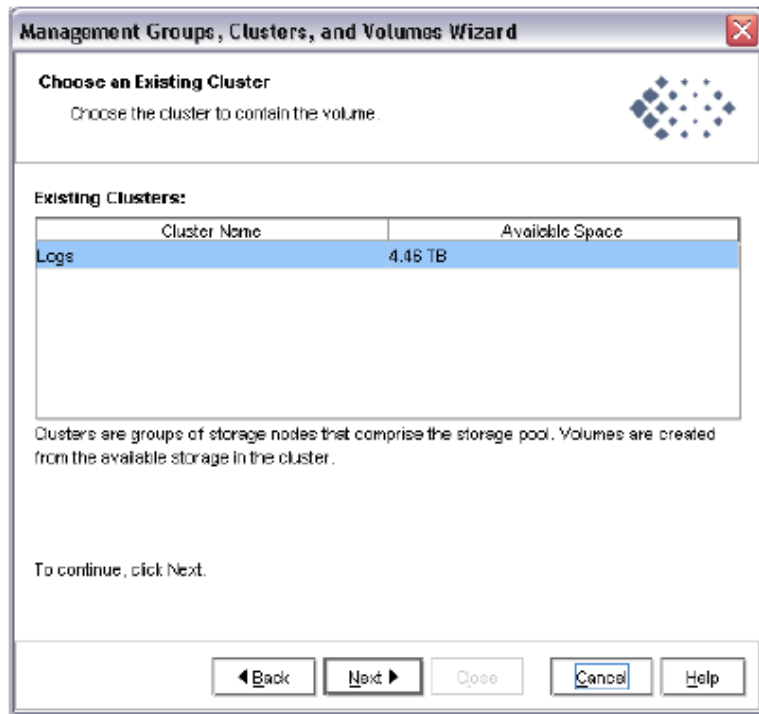


Figure 26 Selecting the cluster to convert

- 7 Select the cluster and click Next.

The Set up Sites window opens, shown in [Figure 27](#).

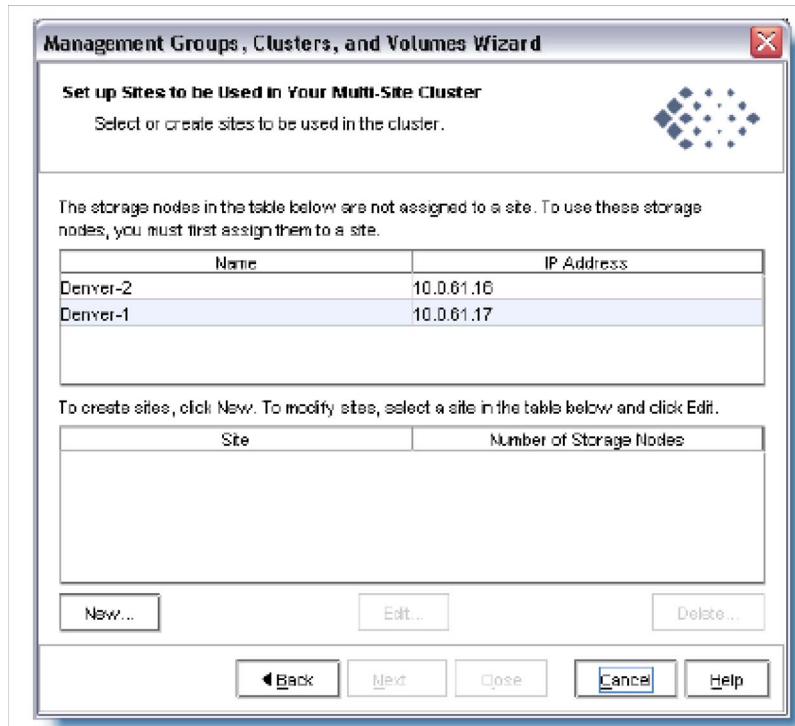


Figure 27 Setting up sites

- 8 Follow the steps found in [Creating Sites](#) on page 40 to create the sites using the listed storage nodes.
- 9 After creating the sites, click Next and the wizard opens the Assign Virtual IPs window, shown in [Figure 28](#). If the current cluster has a VIP assigned, it shows in this window. Use it or assign new ones, according to the network plan for the Multi-Site cluster.

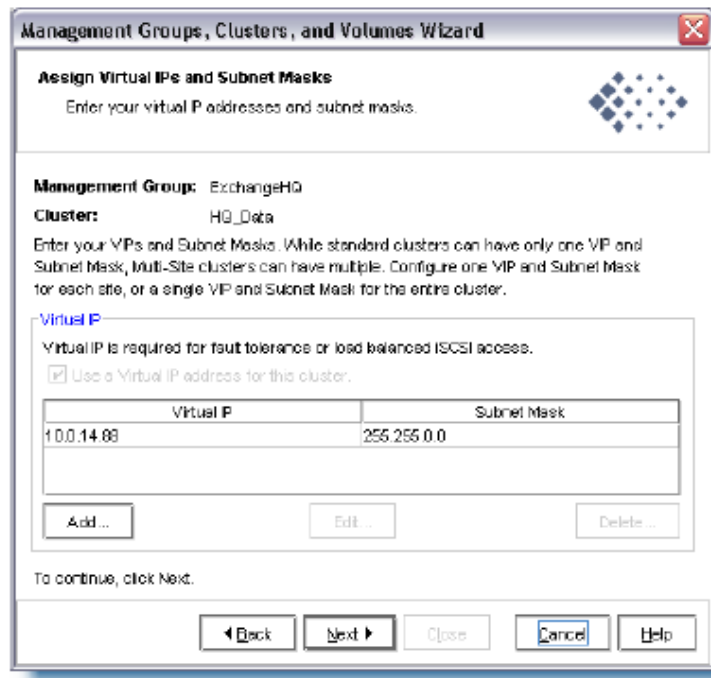


Figure 28 Existing VIP in a standard cluster you are converting to a Multi-Site cluster

- 10 Assign one or more VIPs as planned for your site layout, and click Next. The Create Volume window opens with the box checked to Skip Volume Creation, shown in [Figure 29](#). If you want a new volume for your Multi-Site cluster, clear the box and fill in the volume information.

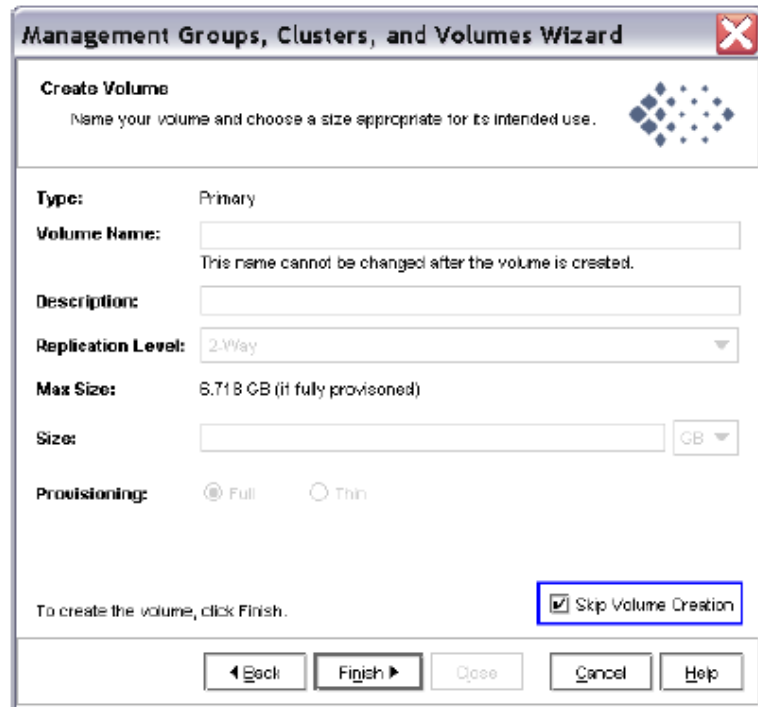


Figure 29 Skip creating a volume if your cluster already has the volumes you want

- 11 Click Next to complete the wizard.
- 12 The Summary window opens where you can verify the settings you have just configured.

Notice in [Figure 30](#) where the newly created sites are listed.

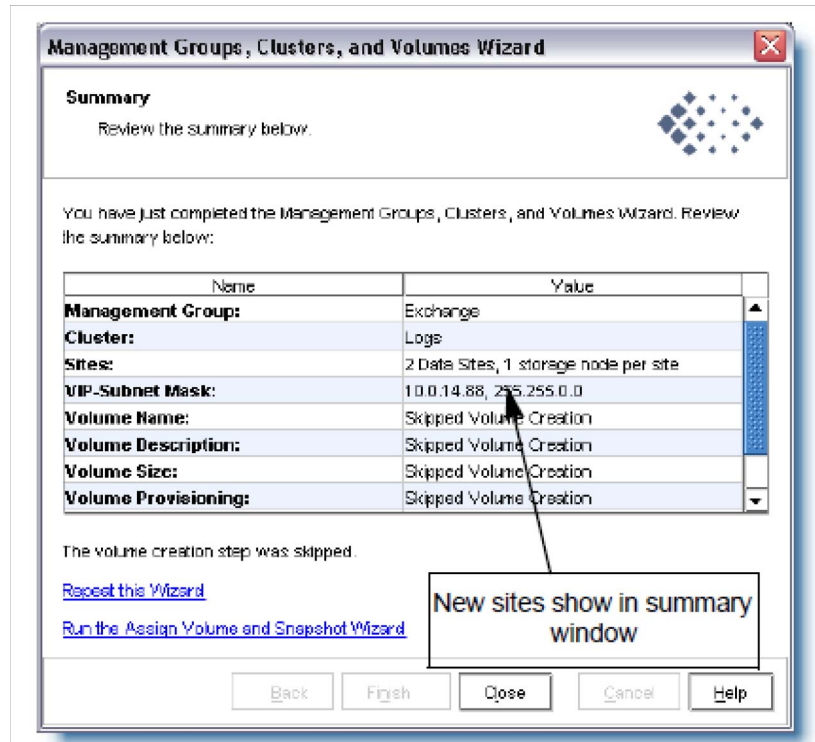


Figure 30 Summary window showing the newly created sites

- 13 Click Close when you finish reviewing the Summary information.
- 14 Select the Sites node in the navigation window, shown in [Figure 31](#), to see the new site configuration.

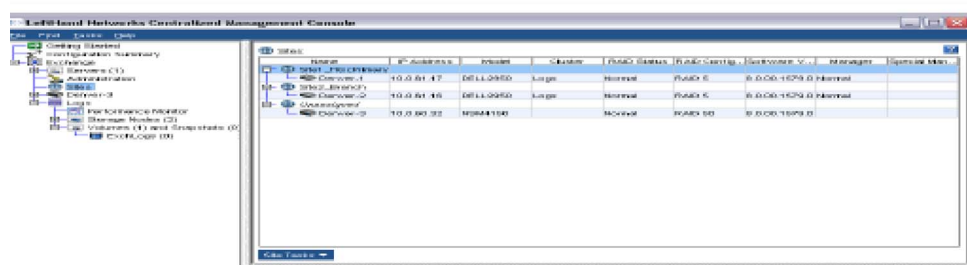


Figure 31 Sites Details table

Configuring Managers

Once you have set up your management group, cluster, and sites, you may need to change the manager configuration of your Multi-Site SAN. You also need to add the Failover Manager in all but one type of configuration.

Managers Overview

The SAN/iQ software automatically starts regular managers on storage nodes when creating a management group. The number of managers started depends on the number of storage nodes in the group and whether you create a Multi-Site cluster. In some configurations, a Virtual Manager is also added to the management group. If you plan to use a Failover Manager, you must manually install the Failover Manager on a separate server on the network and then add it to the management group as required by the design of your Multi-Site SAN. We recommend using the Failover Manager instead of the Virtual Manager because of its automated failover/failback capabilities.

Configuring Regular Managers

After you finish creating the management group, cluster and sites for your Multi-Site SAN, review the manager configuration and change the configuration as necessary. See [Table 4, Common Configurations of Multi-Site SANs and Managers](#) on page 15.

For information about management groups, managers, and starting and stopping managers, see the *LeftHand SAN User Manual*, Chapter 9, “Working with Management Groups,” the section on “Creating a Management Group and Default Managers” on page 271.

Fault Tolerant Manager Configurations

The SAN/iQ software tracks the configuration of managers to ensure that you have a fault-tolerant configuration. You may encounter messages with warnings or prerequisites you must meet before you can change the managers.

For example, if you create a management group and a Multi-Site cluster using 2 or 4 storage nodes, the management group has regular managers running on each storage node plus a Virtual Manager added to the management group. If

you want to add a Failover Manager or stop a manager on one of the storage nodes, you must first delete the Virtual Manager from the management group. Then you can make the other changes.

Installing and Configuring a Failover Manager

If you are using a Multi-Site SAN configuration that calls for a Failover Manager, you install the Failover Manager from the SAN/iQ Management Software CD onto a separate server on the network. The Failover Manager can run on VMware ESX Server, VMware Server or VMware Player. Detailed instructions for installing and configuring the Failover Manager are in the LeftHand SAN User Manual, Chapter 10, “Using Specialized Managers.”

Best Practice

Install the Failover Manager on a server in a separate physical location than the SAN sites and then add it as a separate site in the management group. Doing so will prevent loss of quorum in a site failure scenario, regardless of which single site fails.

Adding the Failover Manager to the Management Group and Site

Adding the Failover Manager to the management group will likely cause the manager configuration to become less than optimal. Be sure to change the manager configuration to best reflect the new configuration of the management group. See the *LeftHand SAN User Manual*, Chapter 9, “Working with Management Groups,” the section “Optimum Number of Managers.”

- 1 Use Find to search for the Failover Manager in the CMC. It appears in the Available pool, shown in [Figure 32](#).

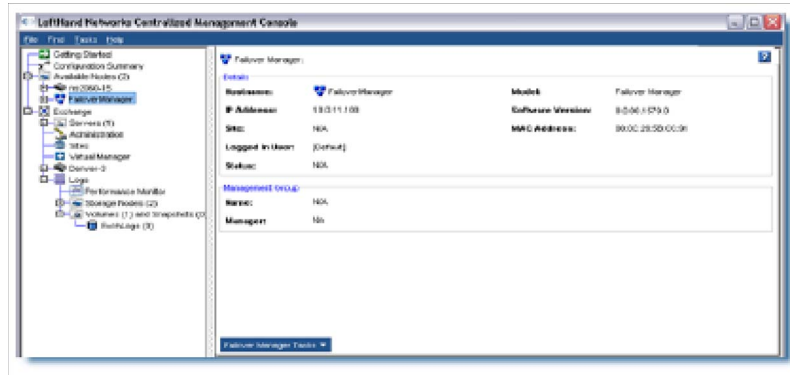


Figure 32 Finding the Failover Manager in the Available pool

- 2 Right-click on the Failover Manager and select Add to Existing Management Group.
 - If there is more than one management group, select the correct one from the list that opens.

If there is only one management group, the Failover Manager is added to the group.

- 3 Select the Sites node in the navigation window.
- 4 Right-click and select New Site.

The New Site window opens.

- 5 Type the name and optional description for the Failover Manager site.
- 6 Click Add in the Site Nodes section.

The Add Nodes to Sites window opens.

- 7 Select the Failover Manager from the list and click OK.
- 8 Click OK on the New Site window.

The new Failover Manager site appears in the Sites Detail tab.

► **Note:** You can also add the Failover Manager to an existing site.

Sites Overview

A site is a user-designated location in which storage nodes are installed. Multi-Site SAN configurations have multiple sites with storage nodes in each site, and each site has its own subnet. A site can be a logical configuration, such as a subnet within the same data center, department, or application. Volumes are replicated across sites, which ensures high availability and redundancy.

Characteristics of sites

- Have equal numbers of storage nodes, except for the Failover Manager, which can reside in a site by itself.
- Correlate to a geographical location and/or a configuration entity such as a subnet or power grid
- May be marked as primary
- Have a planned number of managers
- Have a planned volume replication level

Viewing Sites in the CMC

View sites in the CMC by selecting the Sites node in the navigation window. The Sites node is displayed in every management group. However, typically the Sites Details tab shows the storage nodes as unassigned, shown in [Figure 33](#). The site designation has no effect outside of the Multi-Site SAN configuration.

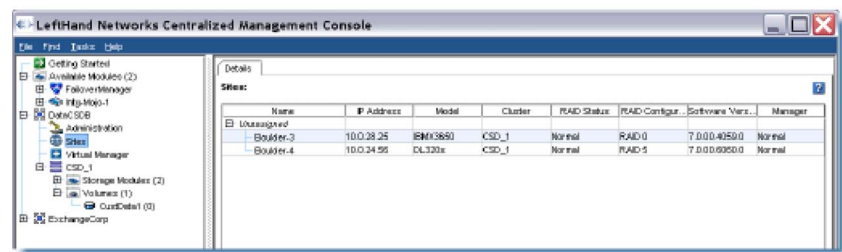


Figure 33 No sites in a standard cluster

When there is a Multi-Site cluster in the management group, the Sites Details tab shows the site assignments of the storage nodes, shown in [Figure 34](#).

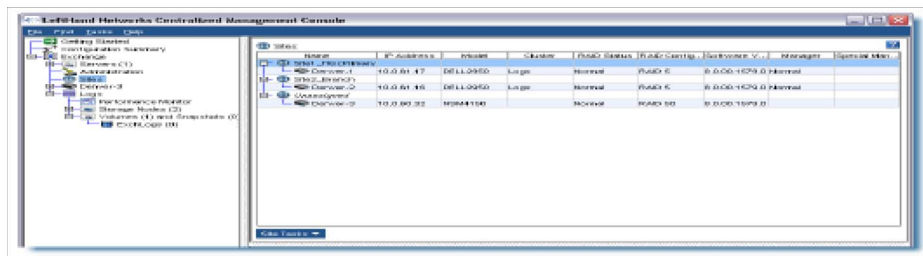


Figure 34 Site assignments in a Multi-Site cluster

A management group can contain multiple clusters, some of which are Multi-Site and some which are not. The Sites tab shows which clusters have sites and which do not. See [Figure 35](#).

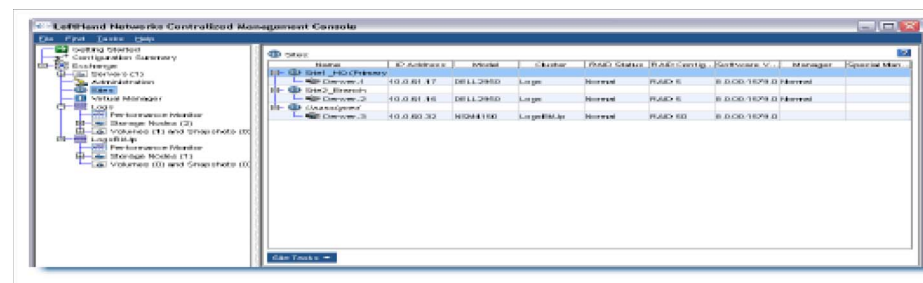


Figure 35 Management groups can have both standard and Multi-Site clusters

Designating a Primary Site

Use the Primary designation when you are using the Multi-Site cluster to maintain a primary site and a second site as a backup site. Designate the site as Primary when you create the site, or you can edit an existing site to make it primary. Only one site can be designated as Primary in a Multi-Site cluster. Make sure that the majority of managers are in the Primary site.

Adding A Site

Create new sites directly in the CMC if you want to add a site to an existing cluster. You must have equal numbers of storage nodes in each site. The Failover Manager can occupy a site by itself.

- 1 Select the Sites node in the navigation window.
- 2 Right-click and select New Site.

The New Site window opens, shown in [Figure 36](#).



Name	IP Address	Model	RAID Conti...	Software ...
------	------------	-------	---------------	--------------

Figure 36 Creating a new site

- 3 Type in a site name and optional description.
- 4 (Optional) Select the box to make the site primary, if that is the configuration you are setting up.
- 5 Click Add to add storage nodes to the site.

The Add Nodes to Site window opens, shown in [Figure 37](#).

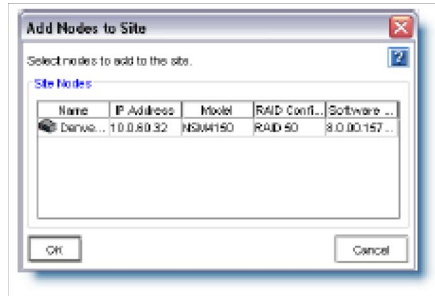


Figure 37 Selecting nodes to add to a new site

6 Select the appropriate nodes and click OK when you are finished.

7 Click OK again when you have finished creating the new site.

The Sites Details tab displays the new site with its storage node(s).

Editing A Site

- You can edit any aspect of the site, including
- Changing the name and description of the site
- Changing the Primary designation (Note: only one site can be primary.)
- Adding or removing nodes. All sites must have the same number of storage nodes, except for the Failover Manager, which can reside in a site by itself.

To Edit a Site

1 Right-click the Sites node in the navigation window and select Edit Sites.

The Site List window opens, shown in [Figure 38](#).

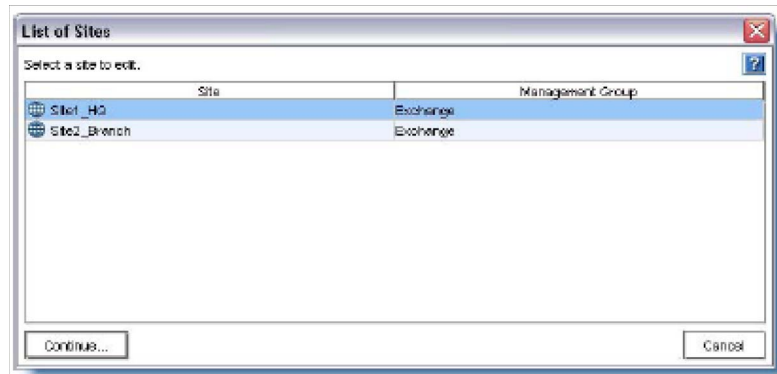


Figure 38 List of sites to edit

- 2 Select the site you want to edit and click Continue.
- 3 The Edit Site window opens, shown in [Figure 39](#).

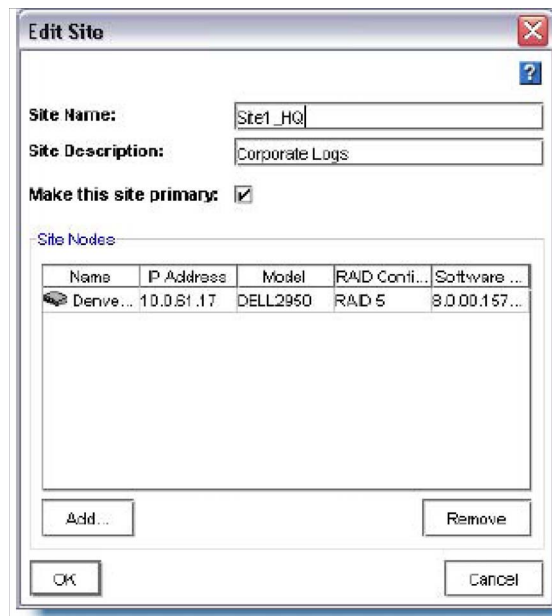


Figure 39 Editing a site

Make the desired changes.

Click OK when you are finished.

Deleting a Site

If you are deleting a site from a Multi-Site cluster, you are either reconfiguring your SAN or recovering from a disaster or other failure.

Deleting a site changes the cluster from a Multi-Site cluster to a standard cluster. If you then remove the storage node(s) assigned to that site, data that resides on the cluster will restripe.

- 1 Select the Sites node in the navigation window.

The Sites Details tab opens.

- 2 Select the site in the Details tab, right-click and select Delete Site.

The warning message opens, shown in [Figure 40](#).

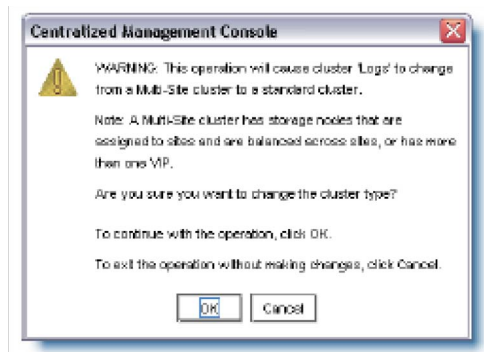


Figure 40 Deleting a site that contains a storage node causes this warning

- 3 Click OK to delete the site.

Adding Storage Nodes to a Multi-Site Cluster

After you have your Multi-Site cluster configured and operating, you may want to add storage nodes to the cluster for additional capacity or to create a new site. While you can add storage nodes to the Multi-Site cluster at any time, you should add them in a way that does not violate the requirements of the Multi-Site SAN.

See [Adding A Site](#) on page 60 if you plan to add the storage nodes to a new site in the cluster.

Adding Additional Capacity to Existing Sites

Plan the following items according to your overall Multi-Site SAN configuration.

- Plan to add the same number of storage nodes to each site.
- Obtain feature keys for the new storage nodes.
- Plan the IP addresses and locations for the new storage nodes to fit into the overall design for your Multi-Site SAN.

To Add Capacity

- 1 Add the license keys to the new storage nodes.
See “[Registering Add-on Applications](#)” on page 531 of the *LeftHand SAN User Manual* for more information about license keys.
- 2 Add the new storage nodes to the management group.
- 3 Select the management group Sites node in the navigation window.
The Sites Details tab opens.
- 4 Select the site to which to add a storage node.
- 5 Right-click and select Edit Site.
The Edit Site window opens, shown in [Figure 41](#).

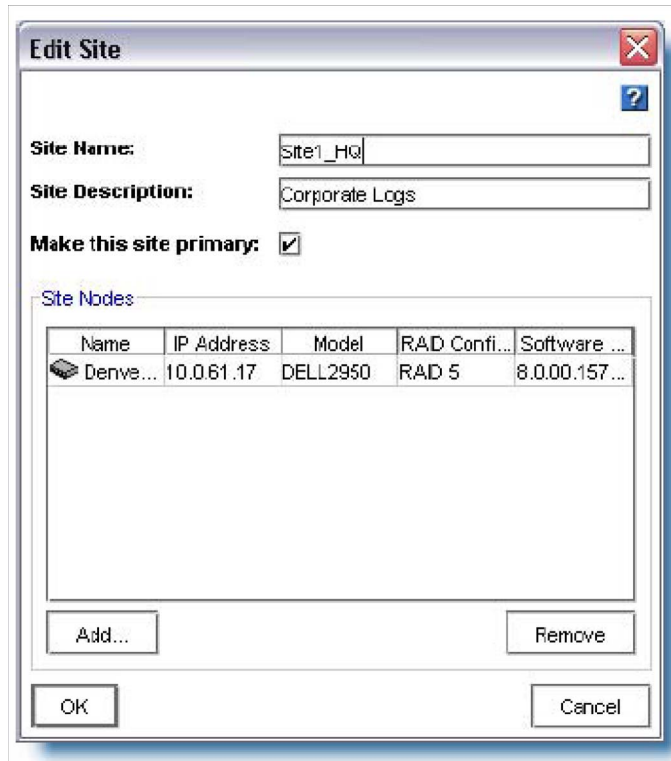


Figure 41 Editing a site to add capacity

- 6 Click Add to add the storage node.
- 7 The Add Nodes to Site window opens.
- 8 Select the storage node from the list and click OK.
- 9 Click OK to finish adding the storage node to the site.
- 10 Repeat [step 4](#) through [step 9](#) for all the storage nodes you are adding.
- 11 Finally, add the storage nodes to the Multi-Site cluster, using the Edit Cluster window.

Removing Storage Nodes from a Multi-Site Cluster

Removing storage nodes from a Multi-Site cluster affects the capacity of the cluster and the Multi-Site configuration.

- Ensure that the capacity remaining after you remove the storage node(s) is sufficient for the data on the volumes.
- To maintain the Multi-Site configuration, plan to remove equal numbers of storage nodes from each site.

To Remove the Storage Nodes from the Site

- 1 Select the Sites node in the navigation window.
The Sites Details tab opens.
- 2 Select the site from which you want to remove the storage node.
- 3 Right-click and select Edit Site.
The Edit Site window opens.
- 4 Select the storage node from the list and click Remove.
The storage node is removed from the list.
- 5 Click OK on the Edit Site window.

A message opens, shown in [Figure 42](#), warning that the Multi-Site cluster will change to a standard cluster if you remove the storage node.

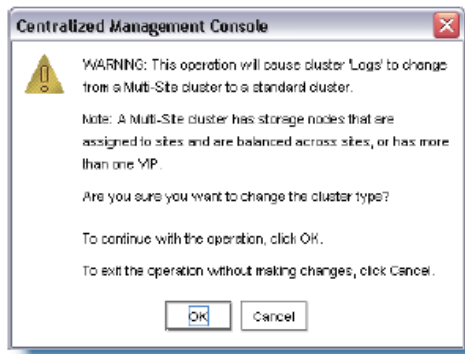



Figure 42 Warning that removing storage node from the Site changes the cluster type

- 6 Click OK to continue.
- 7 Repeat [step 2](#) through [step 5](#) to remove storage nodes from additional sites.

 **Note:** Adding the Failover Manager will likely cause the manager configuration to become less than optimal. Be sure to change the manager configuration to best reflect the new configuration of the management group. See the section “Optimum Number of Managers” in Chapter 9, “Working with Management Groups” of the LeftHand SAN User Manual.

Index

A

- Adding
 - additional capacity to existing sites, 64
 - site, 60
 - storage modules to Multi-Site cluster, 64
- Alternative Multi-Site SAN Configurations, 19
- Application Clusters
 - setting up in a Multi-Site SAN environment, 29

B

- Before You Begin, 35
- Benefits of a Multi-Site SAN, 7
- Best Practices
 - failover manager, 22
 - installing Failover Manager, 56
 - quorum and manager configuration, 31

C

- Capacity, Adding to Sites, 64
- Changing Cluster Type, 63
- Clusters
 - changing from Multi-Site to standard, 63
 - creating with Multi-Site wizard, 37
 - definition, 10
 - removing storage modules from Multi-Site, 66
 - VIP requirement, 10

- Common Multi-Site SAN Network Designs, 23

- Concepts and Terms, 9

- Configuration Diagram
 - dual redundant links between 2 sites and Failover Manager in 3rd site, 25
 - triangular network with Failover Manager, 26

- Configurations, Alternative Multi-Site SAN, 19

- Configuring
 - data replication, 12
 - Failover Manager, 56
 - managers, 55
 - regular managers, 55
 - sites and quorum managers, 15

- Console, Viewing Sites in, 58

- Converting an Existing SAN to a Multi-Site SAN, 46

- Creating
 - management group, 37
 - Multi-Site clusters and volumes, 37

D

- Data Center, Definition as Site, 10
- Data Replication
 - configuring for high availability and redundancy, 12
- Definitions, 9

Deleting a Site, 63

Designating a Primary Site, 59

Designing

Multi-Site SAN, 12

Designing Network for Multi-Site SAN, 23

Disaster Recovery Site definition, 10

DSM for MPIO and Multiple Subnets, 28

E

Editing A Site, 61

F

Failover Manager

best practices, 22

capability of Multi-Site SAN, 7

common Multi-Site SAN configurations

with managers, 15, 33

configuration with dual redundant links

between 2 sites and Failover

Manager in 3rd site, 25

configuration with triangular network

with Failover Manager, 26

configuring, 56

definition, 11

in logical site, 9

overview, 8

typical configuration with 3rd site, 16

G

Glossary, 9

cluster, 10

data center, 10

disaster recovery site, 10

failover manager, 11

logical site, 9

peer site, 9

primary site, 9

secondary site, 9

site, 9

H

High Availability

data replication, 12

site and manager configuration, 15

I

Implementing the Multi-Site SAN, 35

Installation

preparing for, 35

Installing the Multi-Site SAN, 36

L

Logical Site

definition, 9

M

Management Group

creating, 37

Manager

- common configurations with Multi-Site SANs, 15
- configuration best practices, 31
- configuring Failover Manager, 56
- configuring for high availability and redundancy, 15
- configuring regular, 55
- failover, for Multi-Site SAN, 7
- failover, in typical 3rd site configuration, 16
- failover, logical site for, 9
- overview, 55

Multi-Site Cluster

- creating, 37
- removing storage modules from, 66

Multi-Site SAN

- alternative configurations, 19
- capabilities, 7
- common configurations with recommended Failover Manager, 33
- common network designs, 23
- converting from an existing SAN, 46
- designing, 12
- multiple VIPs as feature of, 27
- overview, 7
- preparing to implement, 35
- using multiple VIPs, 22

N

Network

- requirements for performance considerations, 22
- sample configurations for Multi-Site SANs, 24

Network Design

- common designs for Multi-Site SAN, 23

O

Overview

- Failover Manager, 8
- managers, 55
- Multi-Site SAN, 7
- sites, 58

P

Peer Site

- definition, 9

Performance and Network Requirements, 22

Planning

- preparing to install, 35

Primary Site

- definition, 9
- designating, 59

Q

Quorum

- configuration best practices, 31
- recover operation, 30

R

Recover Quorum Operation, 30

Redundancy

- data replication for, 12
- site and manager configuration for, 15

Removing

- storage modules from Multi-Site cluster, 66
- storage modules from s site, 66

Requirements, 8

- standard cluster and VIP, 10

S

Sample Recommended Network

- Configurations for Multi-Site SANs, 24

SAN, Converting Existing to Multi-Site, 46

Secondary Site
definition, 9

Setting Up
application clusters in a Multi-Site SAN
environment, 29
DSM for MPIO with multiple subnets, 28
multiple subnets for a Multi-Site SAN,
27

Sites
adding, 60
adding capacity to, 64
configuring for high availability and
redundancy, 15
data center, 10
definition, 9
deleting, 63
designating primary, 59
disaster recovery, definition of, 10
logical, definition of, 9
overview, 58
peer, definition of, 9
primary, definition of, 9
removing storage modules from, 66
secondary, definition of, 9
viewing in the Console, 58

Storage Modules
adding to sites, 64
removing from a site, 66

Subnets
setting up DSM for MPIO with multiple,
28
setting up multiple for a Multi-Site SAN,
27

U

Using the Multi-Site Cluster Wizard, 37

Using the Primary Site Designation in 2-Site
Configurations, 29

V

Viewing Sites in the Console, 58

VIP

as feature of Multi-Site SAN, 27
requirement for cluster, 10
using multiple for Multi-Site SAN, 22

Volumes

creating with Multi-Site SAN wizard, 37

W

Wizard, Multi-Site Clusters, 37